

LA-UR-21-26760

Approved for public release; distribution is unlimited.

Title: OpenShift on OpenStack Lessons Learned

Author(s): Randles, Timothy C.
Johnson, Steven Lee

Intended for: AppRat project meeting

Issued: 2021-07-14

Disclaimer:

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC for the National Nuclear Security Administration of U.S. Department of Energy under contract 89233218CNA000001. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

OpenShift on OpenStack

Lessons Learned

Steve Johnson, HPC-DES
Tim Randles, HPC-DES

15 July, 2021

Overview

1. Early Trials
2. Dora OpenStack testbed environment
 1. RDO fail
 2. RHOSP success
3. Preparing for Developers
4. OpenShift Virtualization (Have you seen *Inception*?)
5. Observations and Near Term Plans



Early Trials

- Started with Code Ready Containers to get a feel for OpenShift
- Supposedly runs on Windows, MacOS, and Linux.
- Very particular about running on Red Hat Linux or CentOS.
- Was able to get it running on MacOS and CentOS.
- First Observation: this is a GUI wrapped around Kubernetes.
- Later: it's more than that.



Dora OpenStack Testbed Environment

- Cray CS300
 - 2 service nodes
 - 18 blades
 - 24 cores of Intel(R) Xeon(R) CPU E5-2630 v2 @ 2.60GHz (Ivy Bridge)
 - 64GB RAM
 - 1TB SATA HDD
 - 1Gig/10Gig/FDR InfiniBand
 - ~200TB NetApp E2700 JBOD storage
- Vendor Liason Enclave
 - very restricted ingress/egress
 - more on this later...



Dora + RDO == fail

- RDO?
 - Red Hat Distribution of OpenStack
 - community supported Red Hat-based OpenStack distribution
 - think: RDO is to Red Hat OpenStack Distribution as Fedora is to Red Hat Enterprise Linux
- Steve encountered problems deploying OpenShift on RDO
 - Red Hat is not interested in supporting OpenShift on a free-as-in-dollars platform



Early Attempts at Installation on OpenStack

- Explored both IPI (very automated) and UPI (customizable) installations.
- Neither worked.
- Each attempted installation failed at various points.
- Going through an old Linux gateway from an isolated enclave and having to use a web proxy complicated this.
- These attempts provided a good experience for how the configuration files and directories fit together.




Dora + RHOSP == success

- RHOSP?
 - Red Hat OpenStack Platform
 - fully supported for hosting OpenShift
 - NOT cheap
- asked for and acquired evaluation licenses from Red Hat
- installing RHOSP is not a trivial process
 - ~2 weeks (not all day every day for 2 weeks), lots of help from Oak Ridge, one call with a Red Hat engineer



DORA RHOSP Architecture



Successful Installation on Red Hat OpenStack Platform 16

- Obtained evaluation licenses for RHOSP
- This permitted the first successful installation (IPI) of OpenStack 4.7.x
- Was able to spin up simple Kubernetes Deployments
- Eventually encountered storage and networking problems
 - persistent volume storage was too small and was not shared
 - when the available storage filled the service failed
 - another controller node took over for the failed service
 - the next controller didn't have access to the original volumes (not shared!)
 - FIXED by standing up dedicated volume storage service
 - in the future would rather provide NFS-mounted shared storage to all controllers
 - external networking was not accessible
 - quick call with Red Hat engineer identified missed configuration
 - fixed without disrupting existing OpenShift installation



Second Successful Installation

- After RHOSP reconfigure, the IPI install completed in under an hour.
- Cleaner object (Swift) and block (Cinder) storage.
- Improvements to Linux gateway
 - installed OpenStack load balancer service
- This is our current running OpenShift, which has gone through 3 or 4 updates.
- Worked with Red Hat on “Day Two” installation tasks.
- Compliance Operation looks very useful for securing the CoreOS nodes.
- Worked through machinesets to configure CoreOS nodes - also useful.

Early missteps and reconfiguration of OpenStack and OpenShift demonstrated the overall resiliency of the system. We were never forced to “nuke and pave” the entire deployment.



Preparing for Developers: LDAP

- Our isolated enclave prevents authentication connectivity to AD/LDAP.
- Brought up a local containerized OpenLDAP instance for users in OpenShift.
- Needed to create OVN load balancers in OpenStack for each service we wanted to expose.
- OpenStack Floating IP assigned to load balancer after it's running allows us to access services on a known IP address. E.g., LDAP on 389/636.



Preparing for Developers: Git

- Application developers store code and images in institutional Git repo.
- This is inaccessible from our enclave.
- Having a local Git server would help replicate their workflow.
- Tried a few different Helm charts and operators, but settled on a Gitea operator-based installation.
- <https://github.com/jharmison-redhat/gitea-operator>
- Clean install into dev's project space.
- Devs have been working with Red Hat for past 2 weeks.



OpenShift Virtualization

- VMs on top of Containers on top of VMs on top of bare metal.
- Yes, it actually works. Tested with RHEL 8.4 and Windows 2019.
- The use case for this would be running one or more Windows VMs that are part of a service application. Web + middleware + DB, for example.
- This is probably less useful for monolithic apps with which the user directly interacts.
- These are probably best handled with a regular VM or in a Virtual Desktop Infrastructure.



OpenShift Observations

- The “project” concept greatly expands on namespaces, allowing some interesting security and access options.
- Load balancing is non-trivial, but once configured it works.
- The *new-app* and *route* constructs are useful.
- Advanced Cluster Security looks interesting.
- Disconnected (air gapped) mode appears to actually work.
- The OpenShift WebUI can be useful.
- The Red Hat support of the product is quite good.
 - They actually scan their container images before release.
 - A new .z version comes out every two weeks.



Near Term Plans

- Continue to host OpenShift on OpenStack for our Devs to test code.
- Consider installing a second ingress controller for general applications.
- Look at additional platforms on which to run OpenShift:
 - Red Hat Enterprise Virtualization (RHEV - soon to be discontinued)
 - oVirt (free RHEV)
 - VMware
 - Plain libvirt (may not be supported)
 - Bare metal
- How about OpenShift alternatives?
 - Vanilla Kubernetes
 - Rancher Kubernetes
 - HPE Container Platform
 - Others?

